DISTRIBUTED BIOSURVEILLANCE SYSTEMS USING SENSIBLE AGENT TECHNOLOGY TO IMPROVE COORDINATION AND COMMUNICATION AMONG DECISION-MAKERS

K. Suzanne Barber, Karen Fullam, Thomas J. Graser, David C. Han, Joonoo Kim, Dung.Lam, Ryan M. McKay, Jisun Park, Marcelo Vanzin

The Laboratory for Intelligent Processes and Systems
Electrical and Computer Engineering
The University of Texas at Austin
{barber, kfullam, graser, dhan, joonoo, dnlam, rmckay, jisun, vanzin}
@lips.utexas.edu

Protecting the population from chemical-biological attacks and outbreaks of infectious disease is a fundamental goal of health agencies, and early warning is critical for an effective response. However, such biosurveillance activities are inherently challenging due to the complexities involved in coordinating participants; determining the reliability of information; and drawing epidemiological conclusions. By applying Sensible Agent (SA) multi-agent system (MAS) technology to the biosurveillance domain, we can reduce the burden on the TDH epidemiologist by distributing and coordinating decision-making, as well as help the TDH manage the uncertainty of incoming data and understand how that uncertainty impacts resulting epidemiological assessments.

1. INTRODUCTION

Protecting the population from chemical-biological attacks and outbreaks of infectious disease is a fundamental goal of health agencies, and early warning is critical for an effective response. However, such biosurveillance activities are inherently challenging due to the complexities involved in coordinating participants; determining the reliability of information; and drawing epidemiological conclusions.

To support epidemiologists and enable early detection, the Laboratory for Intelligent Processes and Systems is applying its Sensible Agent (SA) multi-agent system (MAS) technology to the biosurveillance domain. Specifically, this agent technology has been demonstrated for biosurveillance in support of the Texas Department of Health (TDH). In the current configuration, all data acquired from hospitals, clinics, laboratories, and other sources is funneled to the TDH epidemiologist, who is the sole decision-maker. The proposed configuration reduces the burden on the TDH epidemiologist by distributing decision-making responsibility (based on disease, region, or other category), and the various decision-makers coordinate to improve efficiency and avoid duplication of effort. The proposed configuration also helps TDH manage the uncertainty of incoming data and understand how that uncertainty impacts resulting epidemiological assessments.

The rest of this paper is structured as follows. The next section introduces the BioChem surveillance domain. Section 3 describes the application architecture for a Sensible Agent-based simulation of that domain. Finally, we close the paper with conclusions.

2. THE BIOCHEM SURVEILLANCE DOMAIN

Protecting the population from chemical-biological attacks and outbreaks of infectious disease is a fundamental goal of government entities such as the Center for Disease Control (CDC), as well as state and local health agencies. Early warning is critical for saving lives and implementing an effective response, including characterizing disease sources, preventing proliferation, and treating patients. However, such biosurveillance activities are inherently challenging due to a number of complications:

1. *Coordinating participants and disseminating information:* Biosurveillance requires coordination between local, state, and federal authorities. Local entities such as treatment facilities must collect

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE **01 OCT 2005**	2. REPORT TYPE **N/A**	3. DATES COVERED -
4. TITLE AND SUBTITLE **Distributed Biosurveillance Systems Using Sensible Agent Technology To Improve Coordination And Communication Among Decision-Makers**		5a. CONTRACT NUMBER 5b. GRANT NUMBER 5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)		5d. PROJECT NUMBER 5e. TASK NUMBER 5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **The Laboratory for Intelligent Processes and Systems Electrical and Computer Engineering The University of Texas at Austin**		8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S) 11. SPONSOR/MONITOR'S REPORT NUMBER(S)
12. DISTRIBUTION/AVAILABILITY STATEMENT **Approved for public release, distribution unlimited**		
13. SUPPLEMENTARY NOTES **See also ADM001851, Proceedings of the 2003 Joint Service Scientific Conference on Chemical & Biological Defense Research, 17-20 November 2003. , The original document contains color images.**		
14. ABSTRACT		
15. SUBJECT TERMS		

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**	**UU**	**7**	

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std Z39-18

information and disseminate to decision-making entities, constrained by communication costs as well as time.

2. *Determining the relevance and reliability of information:* The ability of a treatment facility to offer a diagnosis with high confidence is in part a function of resources and expertise at their disposal. Thus, entities receiving symptom and diagnosis information must determine information relevance qualified by source reliability.
3. *Drawing epidemiological conclusions from symptoms and diagnoses:* Significant expertise is required to analyze symptoms and diagnoses to assess public risk. Such assessments consider an entire region, drawing together information from all contributing sources. Often few epidemiologists are responsible for monitoring a large area, resulting in analysis and communication overload.

In Texas, biosurveillance is conducted by the Texas Department of Health (TDH) in cooperation with local and federal authorities. Resident epidemiologists gather patient symptom and diagnostic information from hospitals, clinics, and other treatment facilities from throughout the state. TDH is currently a single point of contact, and has solicited approaches for other network configurations that allow them to reduce burden and provide more rapid response.

3. Sensible Agent Application Architecture

To support epidemiologists and increase the effectiveness of biosurveillance activities, the Laboratory for Intelligent Processes and Systems at the University of Texas at Austin (UT:LIPS) is applying its Sensible Agent (SA) multi-agent system (MAS) technology to the biosurveillance domain.

UT:LIPS has suggested a variety of options for distributing decision-making and information monitoring using agent technology. Figure 1 depicts one configuration of agent roles suggested by UT:LIPS. The existing two-layer network is expanded to three layers, with an interim layer consisting of Sensible Agents acting as "Disease Agents," each dedicated to the analysis of a particular disease being monitored. Symptom and diagnosis data is sent from information sources (e.g., hospitals and clinics) to disease agents, which evaluate the information and send appropriate epidemiological assessments to a centralized TDH agent. These assessments are in the form of an estimated epidemic level ranging from "None" to "Severe". In making an assessment, disease agents incorporate source reliability and the confidence values associated with information assigned by those sources. Disease agents may also collaborate in acquiring and evaluating symptom and diagnosis data.

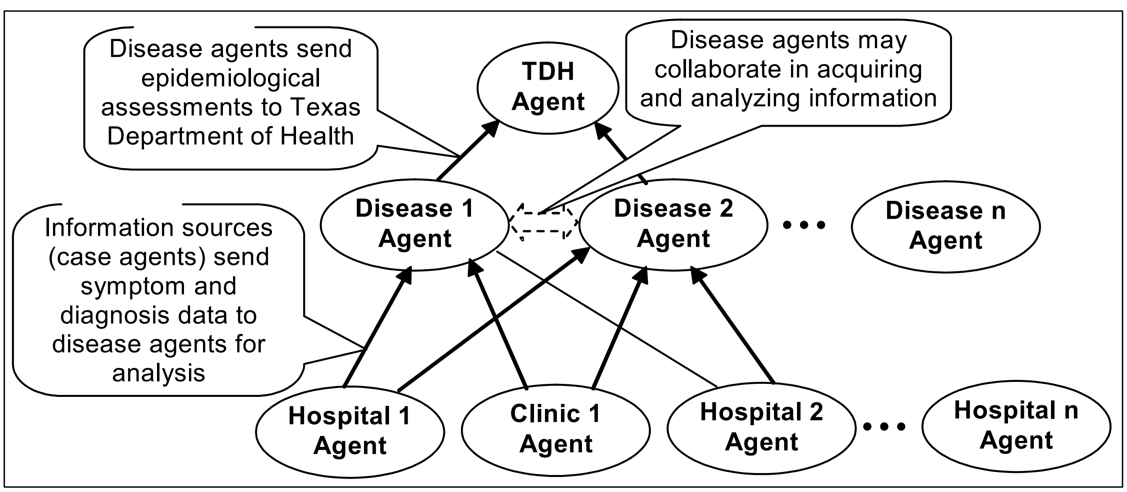

Figure 1: Network of biosurveillance agents.

This paper illustrates the application of Sensible Agents to the biosurveillance domain through various scenarios, each emphasizing one of the following Sensible Agent features: (i) *Evaluation of Information and Information Source Trustworthiness* by considering level of uncertainty of symptom and diagnosis information and reliability of sources transmitting that information to epidemiological decision-makers; (ii) *Adaptive Decision-Making Organization Formation (ADMF)* to establish collaborations among decision-makers when determining epidemiological threat level; and (iii) *Coordinated Planning and Execution* among distributed epidemiological decision-makers, allowing them to share preferences during data acquisition, thereby avoiding extraneous queries.

The application is implemented in the Sensible Agent Testbed. The Testbed offers a suite of tools to assist in the experimental testing of multi-agent systems. Tool facilities include execution visualization and configuration of repeatable experimental runs, allowing the user to specify what parameters to test and what data to record [1].

3.1. EVALUATION OF INFORMATION AND SOURCE TRUSTWORTHINESS

An agent does not have access to ground truth information. Rather, it must build its own local, subjective model of its environment by maintaining beliefs. Beliefs are derived from perceptions sensed or received from information sources via communication. These perceptions have inherent uncertainty; sensors and information sources have accuracy limitations, and some sources may lie about the information they communicate. An agent's belief revision process must assess the quality of incoming information and create a consistent set of beliefs. This research addresses methods for assessing the trustworthiness of information and information sources to produce accurate beliefs and isolate malicious or faulty sources from influence.

The agent's belief revision process is shown in Figure 2. The agent assesses the certainty of source information based on the certainty conveyed by the source and the source's reputation. Current reputation management research [4] focuses on two methods for modeling an information source's trustworthiness. In Direct Trust Revision, a source's reputation is revised based on its past transaction history with the agent. Dissimilarity metrics measure the quality of information received and credit a source's reputation if the source provides accurate information. To conduct Recommended Trust Revision, an agent acquires reputation information from other sources to form its own trust model. In this case, a source's reputation is affected by trust information recommended by other agents. This trust evaluation is developed in both discrete and continuous domains [5], each of which requires not only special considerations for representing beliefs and trust information, but also unique metrics for assessing information quality. Reputation management strategies provide soft security, attaining high truth accuracy in domains with unreliable or malicious information sources. These sources of fraudulent data can be identified and isolated.

In a proposed biosurveillance scenario, a single Disease Agent monitors two diseases: "Anthrax" and "Bad Beer". Supposed disease cases and their certainties are provided to the Disease Agent by information sources at hospitals, clinics, and laboratories. The disease agent weighs this information based on the reputations of the information sources and the certainty conveyed by the source about the information. Reputations of information sources are revised over time, based on the Direct and Recommended Trust Revision strategies previously described. Revision of information source reputations enables the Disease Agent to accurately assess the quality of information conveyed by information sources of varying competencies. In turn, the Disease Agent can build a more accurate picture of existing disease cases and their locations to enable accurate detection of epidemic levels.

Figure 3 shows the result of a simulation run through the scenario described above. The Anthrax agent's beliefs are displayed. The circles signify geographical locations of cases, with a larger circle indicating more cases in that location. To aid the human user understand the situation, the estimated number of cases is mapped to one of six epidemic levels. Each level has a unique color assigned to it,

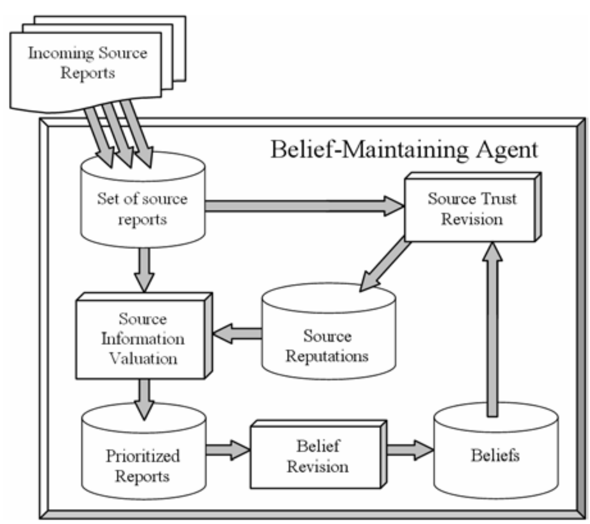

Figure 2: Belief revision process with trust modeling.

and the background of the state of Texas is colored according to the agent's estimated epidemic level In this case, the estimated epidemic level is "HIGH".

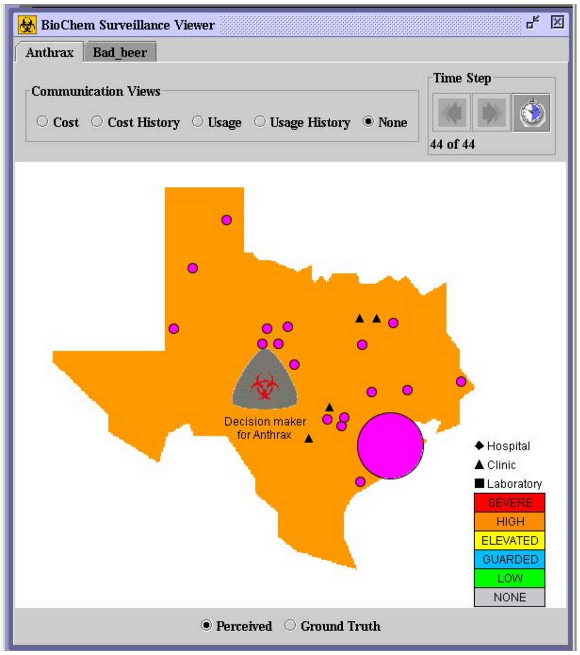

Figure 3: Epidemic evaluation based on information trustworthiness.

3.2. Adaptive Decision-Making Organization Formation

Disease agents will form a Decision-Making Framework (DMF) for each of their respective goals: (1) Determine estimated epidemic level, (2) Obtain symptoms and diagnosis and (3) Evaluate symptoms and diagnosis. Disease agents will employ *Adaptive* DMF (ADMF) capabilities to search for potential partner agents, evaluate those partners and rank their capabilities, select one or more partners, and determine the "best" distribution of decision-making control and execution obligations among selected partners [2].

A DMF specifies how agents work together to decide and execute a given set of goals. A DMF representation has been previously defined as an assignment of variables in three sets, ({D}, {C}, {G}) [2]. This Decision-Making Framework (DMF) representation models the set of agents {D} deciding a set of goals for another, controlled, set of agents {C}, which are bound to accept sub-goals to accomplish the goal set {G}. Agents form a DMF for one or more goals and an agent may participate in multiple DMFs for different goals simultaneously. A Global Decision-Making Framework (GDMF) is a partition of the system's goals and agents set into DMFs so that, at any time, each goal is in exactly one DMF.

A Decision-Making Framework (DMF) is composed of a coherent set of individual decision-making styles for all participating agents (e.g. a Master/Command-driven framework, an All Consensus framework, etc.). Three discrete categories of decision-making interaction styles define salient points along a spectrum of decision-making interaction styles Figure 4 [2].

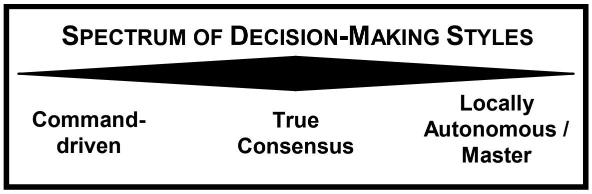

Figure 4: Decision-making interaction styles.

- **Command-driven (CD) –** The agent does not make any decisions about how to pursue its goal and must obey orders given by some other agent(s).

- **True Consensus (CN) –** The agent works as a team member, sharing decision-making control equally with all other decision-making agents.
- **Locally Autonomous (LA)/ Master (M) –** The agent makes decisions alone and may or may not give orders to other agents.

Multi-agent systems capable of Adaptive Decision-Making Frameworks (ADMF) have the ability to change, at run-time, the individual decision-making interaction style of each agent with respect to each goal and, thus, have the ability to make run-time changes to the Decision-Making Frameworks (DMF) in which each agent participates. Multi-agent systems that are not capable of ADMF must use *static decision-making frameworks*, which are established prior to system start-up and employed throughout system operation.

In the biosurveillence scenario, we have two disease agents (Agent 1, Agent 2) and a human agent at the Texas Department of Health (TDH). Agents 1 and 2 are responsible for monitoring "Anthrax" and "Bad Beer" respectively. There is a system policy in place which states that if the estimated epidemic level of any disease is equal to or higher than "Elevated", the human agent can take master control over the Sensible Agent monitoring that disease. Adaptive DMF has been demonstrated in this setup. Figure 5 graphically depicts the progression of decision-making frameworks through this scenario.

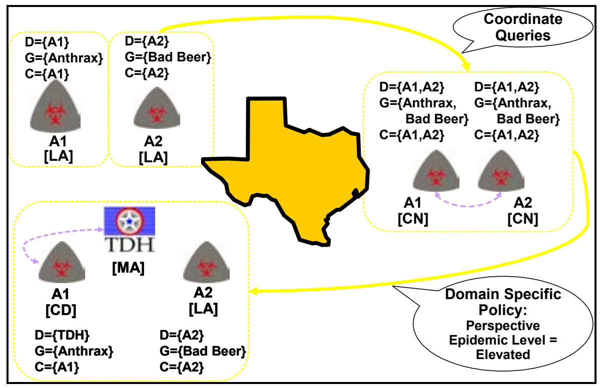

Figure 5: Biosurveillance scenario focusing on ADMF.

Initially, the disease agents work in LA (Locally Autonomous) mode. Meanwhile, both agents are internally exploring the space of feasible DMFs in search of an improvement. In this case, one of the agents decided that a consensus-style DMF with the other agent would be desirable (to coordinate queries). Because the new DMF involves more than one agent, all the involved agents must agree to participate. Once they have, the DMF is formed. Later in the scenario, the epidemic level for Anthrax reaches "Elevated", which triggers the domain-specific system policy mentioned above. Therefore, the human agent becomes the master of a new Anthrax-monitoring DMF, with Agent 1 being command-driven. Agent 2 reverts to Locally Autonomous and monitors Bad Beer by itself.

3.3. Coordinated Planning and Execution

In order to satisfy their goal of maintaining a level of knowledge about the current status of a disease in Texas, the corresponding disease agent must select some set of information sources to query at the current time.

Action selection in this implementation takes advantage of regularities in the domain to simplify the reasoning process. Since each disease agent holds a single goal, the value of a state is equal to the value gained from that single goal (i.e., $V(s) = V_g(s)$). **Equation 1** shows the modified value function used in this domain. The base case provides some reward if the amount of knowledge held in a state, $K(s)$, is greater or equal to some threshold specified in the goal, K_g. The agents represent the query cost by incorporating it into the discount rate, β, as adapted from the goal-based value function defined in equation (3). Instead of defining β as a constant, it is defined as a function, $\beta : A \mapsto [0,1]$, mapping actions into the range of real numbers between 0 and 1.

In addition to the goal, the desires of the disease agent are also manifested in a preference ordering

$$V(s) = \begin{cases} R & \text{if } K(s) \geq K_g \\ \max\limits_{a \in A} \left\{ \beta(a) \cdot V_g(T(s,a)) \right\} & \text{otherwise} \end{cases}$$

Equation 1

on the information sources. This preference ordering represents the value the disease agent places on prospective information that can be retrieved from each source. This preference ordering can be considered to accommodate such features as the amount of time since the information sources were last queried, the historical usefulness of the data provided by information sources, and geographical data about the information sources that handles transmission characteristics of the monitored disease.

Initially, $\beta(a)$ is defined for the agents based on their preference ordering. If action a_i is preferred over action a_j, then $\beta(a_i) > \beta(a_j)$. Given no other outside knowledge, a disease agent will choose to query its most preferred information sources until it obtains the requisite amount of knowledge specified by its goal.

An estimate of the cost of querying an information source is another feature that influences the calculation of $\beta(a)$. To increase the efficiency of the system, the disease agents can communicate some portion of their desires, namely their preference orderings on information sources, so that they may individually make better decisions about which information sources to query. This information is incorporated into decision-making by decreasing $\beta(a_i)$ if a_i is highly preferred by other agents. After recalculating their value functions using the new $\beta(a)$ values, the query load is better balanced across the information sources than if the agents' original value functions were used.

Figure 6 shows a screen capture of the Biosurveillance domain simulation. The lines indicate the costs incurred by querying some set of information sources (marked on the map), with thicker lines indicating higher costs. Coordination among disease agents is intended to spread the information gathering load across the information sources, minimizing the overall cost incurred by the system while still maintaining a reasonably accurate model of the diseases.

The quality of load balancing was used as a comparison metric for the performance of the disease agents when they did or did not share their preference models. Figure 7 shows the experimental results for runs of two homogeneous disease agents and 26 case agents (information sources). Sampling rate is a value defined in the disease agents' goals describing how much information (as a percentage of the number of total information sources) each disease agent must gather for its goal to be considered satisfied. The system cost is a summation of the individual costs to each information source, as a system-wide measure of the load. The cost to an individual information source i at time t is given as a nonlinear function of the number of queries it received in that time step: $Cost_i(t) = queries_i(t)^2$. Optimal is calculated assuming the minimum possible overlap of queries by the disease agents. According to the results, communication of desires (preferences) does indeed improve the system performance.

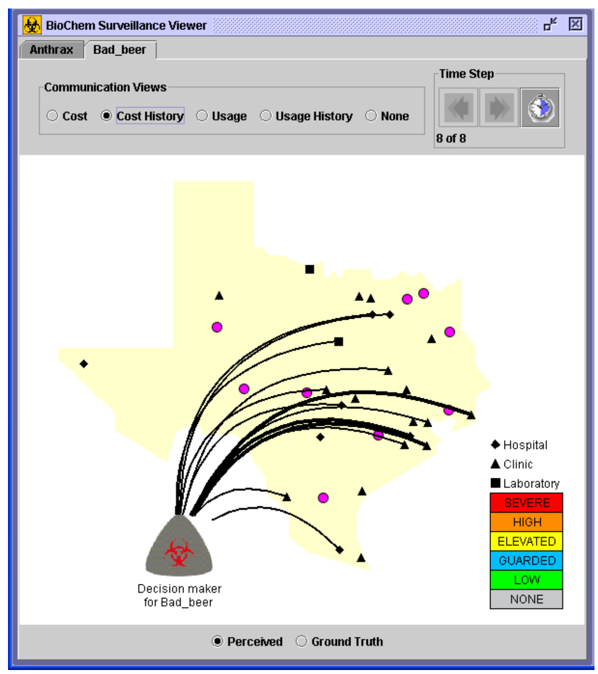

Figure 6: Visualizing query costs.

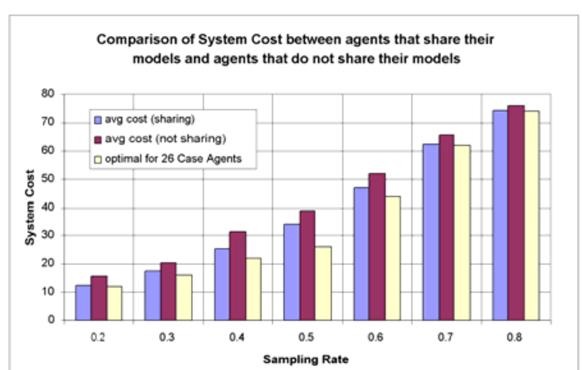

Figure 7: Coordination Experimental Results.

4. CONCLUSIONS

Sensible Agents offer a promising solution for biosurveillance. Particularly, Sensible Agents give decision-makers the ability to (1) coordinate their decisions and actions, (2) effectively and objectively determine the trustworthiness of incoming information and information sources, and (3) dynamically form organizations allocating decision-making control and execution authority based on situational conditions and policy. The Sensible Agent solution distributes the disease monitoring problem to reduce complexity and leverage the specialization and adaptation made possible by both distribution and intelligence. Of course, the human decision-maker remains the final authority. Sensible Agents offer collaborative decision-making insuring humans and agents use the right information from the right sources, to proficiently and adaptively organize, coordinate, decide and act.

5. ACKNOWLEDGMENTS

This research is sponsored in part by the Defense Advanced Research Project Agency (DARPA) Taskable Agent Software Kit (TASK) program, agreement number F30602-00-2-0588. Additionally, the research reported in this document/presentation was performed in connection with Contract/Cooperative Agreement number DAAD17-01-D-0001 with the U.S. Army Research Laboratory. The views and conclusions contained in this document/presentation are those of the authors and should not be interpreted as presenting the official policies or position, either expressed or implied, of the U.S. Army Research Laboratory or the U.S. Government unless so designated by other authorized documents.

6. REFERENCES

[1] Barber, K. S., McKay, R. M., MacMahon, M. T., Martin, C. E., Lam, D. N., Goel, A., Han, D. C., and Kim, J. 2001. Sensible Agents: An Implemented Multi-Agent System and Testbed. In *Proc. of the Fifth Int. Conference on Autonomous Agents (Agents-2001)*, Montreal, QC, Canada: 92-99.

[2] Barber, K. S. and Martin, C. E. 2001. Dynamic Reorganization of Decision-Making Groups. In *Proc. of the Fifth Int. Conference on Autonomous Agents (Agents-2001)*, Montreal, QC, Canada: 513-520.

[3] Barber, K. S. and Kim, J. "Belief Revision Process based on Trust: Simulation Experiments," Autonomous Agents 2001, Workshop for Fraud, Deception and Trust in Agent Societies, pp. 1-12, May 28-June 1, 2001, Montreal, Canada.

[4] Barber, K. S. and Kim, J. " Belief Revision Process based on Trust: Agent Evaluating Reputation of Information Sources," In Trust in Cyber-societies: Integrating the Human and Artificial Perspectives, vol. 2246, Lecture Notes in Computer Science, Falcone, R., Singh, M., and Tan, Y.-H., Eds.: Springer, 73-82.

[5] Barber, K.S. and Fullam, K. "Applying Reputation Models to Continuous Belief Revision," 6th Workshop on Deception, Fraud, and Trust in Agent Societies, 2nd International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2003), July 14- 19, 2003, Melbourne, Australia.